

"Express Mail" mailing label number:

EV324252165US

METHOD AND SYSTEM FOR OPERATING SYSTEM QUIESCENT STATE

Craig Chaiken
Muhammed K. Jaber
Kendall C. Witte

5 BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates in general to the field of information handling system maintenance, and more particularly to a method and system for creating an operating system quiescent state to perform maintenance on information handling system processing components.

Description of the Related Art

As the value and use of information continues to increase, individuals and businesses seek additional ways to process and store information. One option available to users is information handling systems. An information handling system generally processes, compiles, stores, and/or communicates information or data for business, personal, or other purposes thereby allowing users to take advantage of the value of the information. Because technology and information handling needs and requirements vary between different users or applications, information handling systems may also vary regarding what information is handled, how the information is handled, how much information is processed, stored, or communicated, and how quickly and efficiently the information may be processed, stored, or communicated. The variations in information handling systems allow for information handling systems to be general or configured for a specific user or specific use such as financial transaction processing, airline reservations, enterprise data storage, or global communications. In addition, information handling systems may include a variety of hardware and software components that may be configured to process, store, and

communicate information and may include one or more computer systems, data storage systems, and networking systems.

Information handling systems have grown steadily in complexity both in the number and types of hardware components used to build information handling systems and in the software components that rely on these hardware components to perform desired functions. Although the increased complexity generally provides improved performance and greater functionality, it also tends to involve increased maintenance and diagnostics to ensure that components operate properly together. For instance, common maintenance upkeep tasks include performing diagnostics periodically or when problems arise, flashing read only memory (ROM), presenting warning screens by the basic input/output system (BIOS) and backing up the system hard disc drive. Increased hardware component complexity also tends to result in increased demands on the operating system that coordinates interaction between hardware components. The WINDOWS operating system has evolved over time to manage the operation of hardware components including tracking the configuration of the information handling system to coordinate communication between the components.

One difficulty that arises during maintenance and diagnostics of information handling system components is that maintenance and diagnostics operations affects operating system interaction with system components as well as interaction between components themselves. For instance, ROM flashing, performance of certain diagnostics and hard disk drive backup generally require that the operating system be placed in a quiescent state in which configuration information maintained by the operating system is held consistent. As another example, display of BIOS warning screens generally requires an operating system quiescent state so that BIOS assertion of control over the display from the operating system driver does not interrupt the function of critical driver code. The conventional method for placing an operating system into a quiescent state is to boot to a secondary operating system that manages loading and execution of maintenance or diagnostics code and then re-boot into the primary operating system to ensure proper configuration of components and complete execution of the maintenance or diagnostics code. For instance, maintenance and diagnostics on an information handling system with the WINDOWS operating system

typically is performed by booting to DOS and configuring the information handling system to execute the maintenance or diagnostics function on the next reboot before loading the WINDOWS operating system. However, placing an operating system into a quiescent state and rebooting to perform maintenance and diagnostics functions
5 generally is a slow process that delays and frustrates information handling system users.

SUMMARY OF THE INVENTION

Therefore a need has arisen for a method and system which places an operating system into a quiescent state for maintenance and diagnostics without a
10 reboot of the information handling system.

In accordance with the present invention, a method and system are provided which substantially reduce the disadvantages and problems associated with previous methods and systems for placing an information handling system into a quiescent state. An operating system power down state creates a quiescent state that supports
15 running of a quiescent function by intercepting the operating system power down command to components of the information handling system. The quiescent function commands the operating system to return to an operation state without requiring a reboot of the information handling system.

More specifically, a quiescent state utility associated with the operating
20 system sets up a quiescent state to run a quiescent state function by setting a flag in the information handling system BIOS to indicate a quiescent state instead of a power down state and by setting a pointer in the BIOS to point to the quiescent state function module. The quiescent state utility then commands the operating system to enter a power down state, such as the ACPI S3 suspend state. The operating system stores
25 component configuration information in memory and issues a power down command to the BIOS. A quiescent state module associated with the BIOS intercepts the power down command if the quiescent state flag is set and references the pointer to initiate the quiescent state function module without allowing power down off the information handling system components. The quiescent state function module performs a
30 function that requires a quiescent state, such as ROM flash, diagnostics, maintenance,

BIOS warning screen or hard disc drive backup, and then indicates recovery from the quiescent state to the BIOS. Recovery is similar to recovery from a reduced power consumption S3 state that does not require a re-boot of the information handling system.

5 The present invention provides a number of important technical advantages. One example of an important technical advantage is that an information handling system operating system is placed into and recovered from a quiescent state without a reboot of the system. The use of a simulated ACPI S3 reduced power consumption state places the operating system in a quiescent state similar to that achieved with a
10 system shutdown while still keeping power to system components to allow performance of maintenance and diagnostics functions or display BIOS warning screens. Recovery from the simulated S3 state is quicker than performing a system reboot and thus reduces the delay imposed on a user by performance of maintenance and diagnostics that require an operating system quiescent state. Further, the ACPI
15 infrastructure is generally available within information handling systems so that implementation of an operating system quiescent state without a system re-boot requires only minor modifications to existing or newly manufactured information handling systems. The operating system is placed in a quiescent state through use of the ACPI infrastructure to safely allow full control of the system for maintenance and
20 diagnostics yet safely return the system to its initial operational state after the maintenance or diagnostics code is placed into operation.

BRIEF DESCRIPTION OF THE DRAWINGS

25 The present invention may be better understood, and its numerous objects, features and advantages made apparent to those skilled in the art by referencing the accompanying drawings. The use of the same reference number throughout the several figures designates a like or similar element.

Figure 1 depicts a block diagram of an information handling system configured to enter an operating system quiescent state; and

Figure 2 depicts a method for performing information handling system maintenance and diagnostics in an operating system quiescent state.

DETAILED DESCRIPTION

Information handling system functions that require an operating system quiescent state are performed without a system reboot by inducing an operating system quiescent state with a power down state command and intercepting operating system power down commands to components of the information handling system. For purposes of this application, an information handling system may include any instrumentality or aggregate of instrumentalities operable to compute, classify, process, transmit, receive, retrieve, originate, switch, store, display, manifest, detect, record, reproduce, handle, or utilize any form of information, intelligence, or data for business, scientific, control, or other purposes. For example, an information handling system may be a personal computer, a network storage device, or any other suitable device and may vary in size, shape, performance, functionality, and price. The information handling system may include random access memory (RAM), one or more processing resources such as a central processing unit (CPU) or hardware or software control logic, ROM, and/or other types of nonvolatile memory. Additional components of the information handling system may include one or more disk drives, one or more network ports for communicating with external devices as well as various input and output (I/O) devices, such as a keyboard, a mouse, and a video display. The information handling system may also include one or more buses operable to transmit communications between the various hardware components.

Referring now to Figure 1, a block diagram depicts an information handling system 10 configured to support execution of a quiescent state function without a system reboot. Information handling system 10 processes information with plural components, including a CPU 12, BIOS 14, RAM 16, ROM 18, hard disc drive 20, south bridge 22, video card 24 and chip set 26. The processing of information by the processing components is coordinated by an operating system 28, such as WINDOWS, which operates on CPU 12 with defined RAM 16 and communicates with processing components through BIOS 14. A quiescent function module 30

performs one or more functions that involve interaction with processing components and require operating system 28 to be in a quiescent state so that operating system 28 will not interfere with the quiescent function or hang-up due to configuration changes made to the processing components by the quiescent function. Quiescent functions
5 include ROM flashing, certain diagnostics, BIOS warning screens and hard disc drive back-up. In a quiescent state, operating system 28 does not exert control over processing components and has its configuration settings saved in memory, such as hard disc drive 20. Quiescent state function module includes an operating system, such as a DOS or Linux kernel, that supports operations by quiescent function code.

10 Quiescent function module 30 initiates its operation by communication with a quiescent state utility 32 of operating system 28. Quiescent state utility 32 provides quiescent function information to a quiescent state module 34 of BIOS 14. The information includes a flag, FAKE_S3_FLAG, to indicate that a quiescent state is imminent and a pointer, FAKE_S3_PTR, that points to the location of quiescent
15 function module so that BIOS 14 will be able to call quiescent function module once operating system 28 enters a quiescent state. Quiescent state utility 32 then commands operating system 28 to enter a power down state, such as the ACPI S3 suspend state. The S3 state is ordinarily used to place an information handling system into a suspend mode in which the operating system configuration is saved in RAM 16
20 and the majority of the remaining processing components are powered down by a D3 signal sent from BIOS 14 or and ACPI controller 36 typically located in south bridge 22. The S3 suspend state conserves power and also allows for a rapid recovery of information handling system 10 from the power down state by reinstituting operating system 28 with information from RAM 16 instead of through a re-boot.

25 On receipt of the S3 suspend command from quiescent state utility 32, operating system 28 saves configuration information to enter quiescence and issues a S3 suspend call to ACPI controller 36. Quiescent state module 34 has an SMM I/O trap that intercepts the suspend call to prevent power down of the components if the quiescent state flag is set. BIOS 14 then calls quiescent function module 30 at the
30 location indicated in the quiescent state pointer and quiescent function module initiates its operations with an operating system kernel to perform the quiescent function. Code within quiescent function module 30 configures components needed

to perform the function or, alternatively, BIOS 14 restores end of post configurations before calling quiescent function module 30. Memory to perform the quiescent function may be reserved by operating system 28 or also set by POST. Upon completion of the quiescent function, quiescent function module 30 signals BIOS 14 to exit the simulated S3 suspend state. In some instances, such as when the quiescent function includes diagnostics that reconfigures the CPU or chipset, BIOS 14 may execute an S3 Resume. The simulated S3 state does not require a reboot of information handling system 10 and thus allows a less intrusive, quicker and quieter quiescent function operation. In addition, with WINDOWS the use of a simulated quiescent state supports the running of legacy DOS utilities and thus reduces the cost of rewriting the utilities for operating on WINDOWS.

Referring now to Figure 2, a flow diagram depicts a process for performing a function with operating system quiescence that avoids a reboot of the information handling system. The process begins at step 38 with setting of the FAKE_S3 BIOS flag to indicate that the next S3 suspend call received at the BIOS is a fake suspend call. The process continues to step 40 at which a WINDOWS utility loads the quiescent function code to be executed once WINDOWS is in a quiescent state. At step 42, the WINDOWS utility stores the location of the quiescent function code in FAKE_S3_PTR in the BIOS. At step 44, the WINDOWS utility issues a power management S3 suspend call for WINDOWS to save configuration information and initiate power down of components in accordance with the ACPI standard. At step 46, the S3 suspend call issued from WINDOWS is intercepted by the BIOS.

At step 48, the BIOS determines if the FAKE_S3 flag is set to indicate that the operating system S3 call was fake. If the flag is not set, the process continues to step 50 for the BIOS to allow the ACPI suspend process continue so that the components of the information handling system enter the suspend state. If the flag is determined to be set at step 48 to indicate that the S3 suspend call is fake, the process continues to step 52 at which the BIOS calls the quiescent function at the location indicated by the pointer FAKE_S3_PTR. At step 54, the quiescent function executes and, at step 56, after completion of the quiescent function the BIOS exits the simulated S3 state to return the information handling system to normal operations.

Although the present invention has been described in detail, it should be understood that various changes, substitutions and alterations can be made hereto without departing from the spirit and scope of the invention as defined by the appended claims.